

CLAIMS

1. A method of determining which of a set of content provider mirror sites should receive an end user's initial content request, comprising:

identifying a set of proxy points, wherein each proxy point represents a given point
5 in the Internet at which a trace originating from each of a set of mirror sites directed toward a given name server intersect;

probing the proxy points to generate given data;

generating a download predictor score for each mirror site based on the given data;

10 identifying which mirror site provides a best download performance based on the download predictor score;

associating a given name server IP address with the identified mirror site; and

in response to an end user's initial content request to a given local name server, returning an IP address of the identified mirror site.

2. A method of optimizing a user's initial request to a content provider web site that is replicated at a set of mirror sites, comprising:

responsive to an end user's local name server making a request to the content provider's web site, directing the request to a global load balancing service having a
5 network map that estimates relative connectivity to the mirror sites from a set of proxy points;

using the network map to return to the end user's local name server an IP address identifying an optimal mirror site at which the request may be serviced.

10 3. The method as described in Claim 2 wherein each core point represents an intersection of trace routes that originate from a mirror site to a local name server.

4. A method of routing a user's initial request to a content provider web site that is replicated at a set of mirror sites, comprising:
responsive to an end user's local name server making a request to the content provider web site, directing the request to a global load balancing service having a network
5 map that estimates relative connectivity to the mirror sites from a set of proxy points;
determining whether the network map includes data associating the end user's local name server to one of the mirror sites; and
if not, identifying a given mirror site to respond to the request using a default routing mechanism.

10

5. The method as described in Claim 4 wherein the default routing mechanism is BGP.

6. The method as described in Claim 4 wherein the default routing mechanism
15 is geo-routing.

7. A method for managing global traffic redirection for a set of content providers operating mirrored sites, comprising:

from each of a set of data centers that host mirrored sites, executing a given network test against each of a set of core points;

5 generating a time-weighted average of a given metric based on data generated by executing the given network test;

generating a score for each data center per core point;

generating a set of candidate data centers for each of a set of name servers;

associating a candidate data center to each of a set of IP address space blocks to

10 generate a map;

providing the map to a name server; and

using the map to direct end user requests to a mirrored site to a given data center.

8. The method as described in Claim 7 wherein the given network test is a
15 ping test.

9. The method as described in Claim 7 wherein the given metric is latency or
packet loss.

10. The method as described in Claim 7 further including the step of discarding
20 from the set of candidate data centers any data center that does not meet a given operating criteria.

11. The method as described in Claim 10 wherein the given operating criteria is
25 evaluated using a file download test.

12. A method of optimizing a client request to a content provider site that is replicated at a set of mirror sites, comprising:

generating a network map that estimates relative connectivity to the mirror sites from a set of proxy points;

responsive to a local name server making a request to the content provider's site, directing the request to a global load balancing service; and

having the global load balancing service use the network map to return to the local name server an IP address identifying an optimal mirror site at which the request may be serviced.

13. The method as described in Claim 12 wherein the client request originates at a client machine and the content provider site is a Web site.

14. The method as described in Claim 12 wherein the client request originates from a cache and the set of mirror sites comprises a plurality of storage servers.

15. The method as described in Claim 12 wherein the client request originates at a streaming server and the set of mirror sites comprises a plurality of signal acquisition points.

16. The method as described in Claim 12 wherein the client request originates at a logging process and the set of mirror sites comprises a plurality of log archival servers.

17. The method as described in Claim 12 wherein the client request originates at a mail process and the set of mirror sites comprises a plurality of mail servers.